How Many People Can Hide in a Terrain?

Stephan Eidenbenz

Institute for Theoretical Computer Science,
ETH Zürich, Switzerland
eidenben@inf.ethz.ch

Abstract. How many people can hide in a given terrain, without any two of them seeing each other? We are interested in finding the precise number and an optimal placement of people to be hidden, given a terrain with n vertices. In this paper, we show that this is not at all easy: The problem of placing a maximum number of hiding people is almost as hard to approximate as the MAXIMUM CLIQUE problem, i.e., it cannot be approximated by any polynomial-time algorithm with an approximation ratio of n^ϵ for some $\epsilon > 0$, unless $P = NP$. This is already true for a simple polygon with holes (instead of a terrain). If we do not allow holes in the polygon, we show that there is a constant $\epsilon > 0$ such that the problem cannot be approximated with an approximation ratio of $1 + \epsilon$.

1 Introduction and Problem Definition

While many of the traditional art gallery problems such as VERTEX GUARD and POINT GUARD deal with the problem of guarding a given polygon with a minimum number of guards, the problem of hiding a maximum number of objects from each other in a given polygon is intellectually appealing as well. When we let the problem instance be a terrain rather than a polygon, we obtain the following background, which is the practical motivation for the theoretical study of our problem: A real estate agency owns a large, uninhabited piece of land in a beautiful area. The agency plans to sell the land in individual pieces to people who would like to have a cabin in the wilderness, which to them means that they do not see any signs of human civilization from their cabins. Specifically, they do not want to see any other cabins. The real estate agency, in order to maximize profit, wants to sell as many pieces of land as possible.

In an abstract version of the problem we are given a terrain which represents the uninhabited piece of land that the real estate agency owns. A *terrain* T is a two-dimensional surface in three-dimensional space, represented as a finite set of vertices in the plane, together with a triangulation of their planar convex hull, and a height value associated with each vertex. By a linear interpolation inbetween the vertices, this representation defines a bivariate continuous function. The corresponding surface in space is also called a 2.5-dimensional terrain. A terrain divides three-dimensional space into two subspaces, i.e. a space above and a space below the terrain, in the obvious way. In the literature, a terrain is also called a *triangulated irregular network* (TIN), see [8]. The problem now consists

A. Aggarwal, C. Pandu Rangan (Eds.): ISAAC'99, LNCS 1741, pp. 184–194, 1999.

of finding a maximum number of lots (of comparatively small size) in the terrain, upon which three-dimensional bounding boxes can be positioned that represent the cabins such that no two points of two different bounding boxes see each other. Two points *see* each other, if the straight line segment connecting the two points does not intersect the space below the terrain. Since the bounding boxes that represent the cabins are small compared to the overall size and elevation changes in the terrain (assume that we have a mountainous terrain), we may consider these bounding boxes to be zero-dimensional, i.e. to be points on the terrain. This problem has other potential applications in animated computer-games, where a player needs to find and collect or destroy as many objects as possible. Not seeing the next object while collecting an object makes the game more interesting. We are now ready to formally define the first problem that we study:

Definition 1. *The problem* MAXIMUM HIDDEN SET ON TERRAIN *asks for a set S of maximum cardinality of points on a given terrain T, such that no two points in S see each other.*

In a variant of the problem, we introduce the additional restriction that these points on the terrain must be vertices of the terrain.

Definition 2. *The problem* MAXIMUM HIDDEN VERTEX SET ON TERRAIN *asks for a set S of maximum cardinality of vertices of a given terrain T, such that no two vertices in S see each other.*

In a more abstract variant of the same problem, we are given a simple polygon with or without holes instead of a terrain. A *simple polygon with holes* in the plane is given by its ordered sequence of vertices on the outer boundary, together with an ordered sequence of vertices for each hole. A *simple polygon without holes* in the plane is simply given by its ordered sequence of vertices on the outer boundary. Again, we can impose the additional restriction that the points to be hidden from each other must be vertices of the polygon. This yields the following four problems.

Definition 3. *The problem* MAXIMUM HIDDEN SET ON POLYGON WITH(OUT) HOLES *asks for a set S of maximum cardinality of points in the interior or on the boundary of a given polygon P, such that no two points in S see each other.*

Definition 4. *The problem* MAXIMUM HIDDEN VERTEX SET ON POLYGON WITH(OUT) HOLES *asks for a set S of maximum cardinality of vertices of a given polygon P, such that no two vertices in S see each other.*

Two points in the polygon see each other, if the straight line segment connecting the two points does not intersect the exterior (and the holes) of the polygon. In this paper, we propose a reduction from MAXIMUM CLIQUE to MAXIMUM HIDDEN SET ON POLYGON WITH HOLES. The same reduction with minor modifications will also work for MAXIMUM HIDDEN SET ON TERRAIN, MAXIMUM HIDDEN VERTEX SET ON POLYGON WITH HOLES, and MAXIMUM HIDDEN

Vertex Set on Terrain. Maximum Clique cannot be approximated by a polynomial-time algorithm with a ratio of $n^{1-\epsilon}$ unless $coR = NP$ and with a ratio of $n^{\frac{1}{2}-\epsilon}$ unless $NP = P$ for any $\epsilon > 0$, where n is the number of vertices in the graph [7]. We will show that our reduction is gap-preserving (a technique proposed in [1]), and thus show inapproximability results for all four problems. Maximum Clique consists of finding a maximum complete subgraph of a given graph G, as usual.

We also propose a reduction from Maximum 5-Occurrence-2-Satisfiability to Maximum Hidden Set on Polygon without Holes, which will also work for Maximum Hidden Vertex Set on Polygon without Holes. Maximum 5-Occurrence-2-Satisfiability is APX-hard, which is equivalent to saying that there exists a constant $\epsilon > 0$ such that no polynomial algorithm can achieve an approximation ratio of $1 + \epsilon$ for Maximum 5-Occurrence-2-Satisfiability. See [3] for an introduction to the class APX and for the relationship between the two classes APX and $MaxSNP$, see [11] for the $MaxSNP$-hardness proof for Maximum 5-Occurrence-2-Satisfiability. Please note that $MaxSNP$-hardness implies APX-hardness [3]. We show that our reduction is gap-preserving and thus establish the APX-hardness of Maximum Hidden (Vertex) Set on Polygon without Holes. Maximum 5-Occurrence-2-Satisfiability consists of finding a truth assignment for the variables of a given boolean formula. The formula consists of disjunctive clauses with at most two literals and each variable appears in at most 5 literals. The truth assigment must satisfy a maximum number of clauses.

There are various problems that deal with terrains. Quite often, these problems have applications in the field of telecommunications, namely in setting up communications networks. There are some upper and lower bound results on the number of guards needed for several kinds of guards to collectively cover all of a given terrain [2]. Very few results on the computational complexity of terrain problems are known. The shortest watchtower (from where a terrain can be seen in its entirety) can be computed in time $O(n \log n)$ [15]. The problem of finding a minimum number of vertices of a terrain such that guards at these vertices see all of the terrain is NP-hard and cannot be approximated with an approximation ratio that is better than logarithmic in the number of vertices of the terrain. Similar results hold for the variation, where guards may only be placed at a certain given height above the terrain [5]. When we deal with polygons rather than terrains, we speak of art gallery or visibility problems. Many results (upper and lower bounds, as well as computational complexity results) are known for visibility problems. See [10,13,14] for an overview, as well as more recent work on the inapproximability of Vertex/Edge/Point Guard on polygons with [4] and without holes [6].

The problems Maximum Hidden Set on a Polygon without Holes and Maximum Hidden Vertex Set on a Polygon without Holes are known to be NP-hard [12]. This immediately implies the NP-hardness of the corresponding problems for polygons with holes. A quite simple reduction from these polygon problems to the terrain problems (as given in Sect. 3) even implies

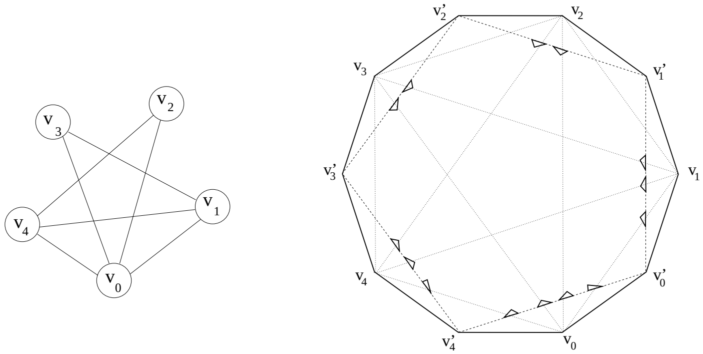

Fig. 1. Example graph and the polygon constructed from it

the NP-hardness for the two terrain problems as well. In this paper, we give the first inapproximability results for these problems. Our results suggest that these problems differ significantly in their approximation properties.

This paper is organized as follows. In Sect. 2, we propose a reduction from Maximum Clique to Maximum Hidden Set on Polygon with Holes. We show that our reduction is gap-preserving and obtain our inapproximability results for Maximum Hidden (Vertex) Set on a Polygon with Holes. We show that our proofs also work for Maximum Hidden (Vertex) Set on Terrain with minor modifications in Sect. 3. In Sect. 4, we show the APX-hardness of Maximum Hidden (Vertex) Set on Polygon without Holes. Finally, we draw some conclusions in Sect. 5.

2 Inapproximability Results for the Problems for Polygons with Holes

Suppose we are given an instance I of Maximum Clique, i.e. an undirected graph $G = (V, E)$, where $V = v_0, \ldots, v_{n-1}$. Let $m := |E|$. We construct an instance I' of Maximum Hidden Set on Polygon with Holes as follows. I' consists of a polygon with holes. The polygon is basically a regular $2n$-gon with holes, but we replace every other vertex by a comb-like structure. Each hole is a small triangle designed to block the view of two combs from each other, whenever the two vertices, to which the combs correspond, are connected by an edge in the graph. Figure 1 shows an example of a graph and the corresponding polygon with holes. (Note that only the solid lines are lines of the polygon and also note that the combs are not shown in Fig. 1.)

Let the regular $2n$-gon consist of vertices $v_0, v_0', \ldots, v_{n-1}, v_{n-1}'$ in counter-clockwise order, to indicate that we map each vertex $v_i \in V$ in the graph to a vertex v_i in the polygon. We need some notation, first. Let $e_{i,j}$ denote the intersection point of the line segment from v_{i-1}' to v_i' with the line segment from v_i to v_j, as indicated in Fig. 2. (Note that we make liberal use of the notation

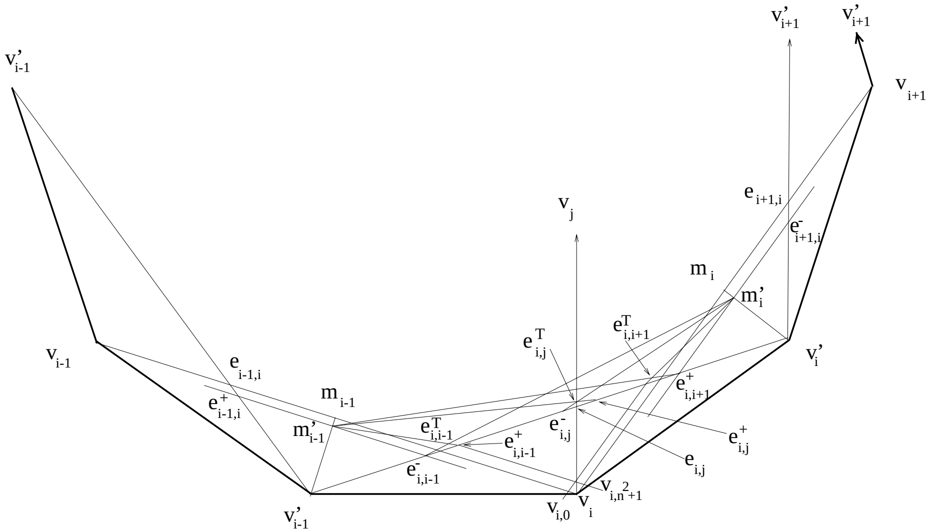

Fig. 2. Points $e^-_{i,j}$, $e^+_{i,j}$, and $e^T_{i,j}$

index for the vertices, i.e. v_{i+1} is strictly speaking $v_{i+1 \bmod n}$, accordingly for v_{i-1}.) Let d denote the minimum of the distances of $e_{i,j}$ from $e_{i,j+1}$, where the minimum is taken over all $i, j = 1, \ldots, n$. Let $e^-_{i,j}$ ($e^+_{i,j}$) denote the point at distance $\frac{d}{3}$ from $e_{i,j}$ on the line from v'_{i-1} to v'_i that is closer to v'_{i-1} (v'_i). Let m_i be the midpoint of the line segment from vertex v_i to v_{i+1} and let m'_i be the intersection point of the line from v'_i to m_i and from $e^+_{i,i+1}$ to $e^-_{i+1,i}$ (see Fig. 2). Finally, let $e^T_{i,j}$ denote the intersection point of the line from $e^-_{i,j}$ to m'_i and the line from $e^+_{i,j}$ to m'_{i-1}. The detailed construction of these points is shown in Fig. 2. We let the triangle formed by the three vertices $e^+_{i,j}$, $e^-_{i,j}$, and $e^T_{i,j}$ be a hole in the polygon iff there exists an edge in G from v_i to v_j. Recall Fig. 1, which gives an example.

We now refine the polygon obtained so far by cutting off a small portion at each vertex v_i. For each $i \in \{0, \ldots, n\}$, we introduce two new vertices $v_{i,0}$ and v_{i,n^2+1} as indicated in Fig. 2. Vertex $v_{i,0}$ is defined as the intersection point of the line that is parallel to the line from v_{i-1} to v_i and goes through point $e^+_{i,i-1}$ and of the line from v_i to v'_i. Symmetrically, vertex v_{i,n^2+1} is defined as the intersection point of the line that is parallel to the line from v_{i+1} to v_i and that goes through point $e^-_{i,i+1}$ and of the line from v_i to v'_{i-1}. We fix n^2-1 additional vertices $v_{i,1}, \ldots, v_{i,n^2}$ on the line segment from $v_{i,0}$ to v_{i,n^2+1} for each i as shown in Fig. 3. For a fixed i, the two vertices $v_{i,l}$ and $v_{i,l+1}$ have equal distance for all $l \in \{0, \ldots, n^2\}$. Finally, we fix n^2 additional vertices $w_{i,l}$ for $l \in \{0, \ldots, n^2\}$ for each i. Vertex $w_{i,l}$ is defined as the intersection point of the line from vertex v'_{i-1} through $v_{i,l}$ with the line from vertex vertex v'_i through $v_{i,l+1}$. The polygon between two vertices v'_{i-1} and v'_i is now given by the following ordered sequence

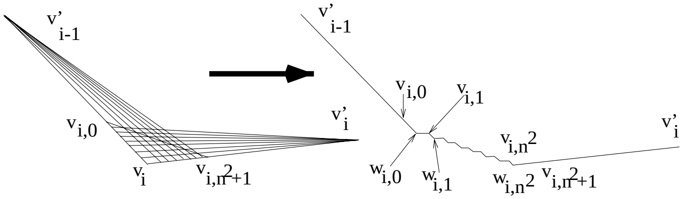

Fig. 3. Construction of the comb of v_i

of vertices: $v'_{i-1}, v_{i,0}, w_{i,0}, v_{i,1}, w_{i,1}, \ldots, v_{i,n^2}, w_{i,n^2}, v_{i,n^2+1}, v'_i$ as indicated in Fig. 3. We call the set of all triangles $v_{i,l}, w_{i,l}, v_{i,l+1}$ for a fixed i and all $l \in \{0, \ldots, n^2\}$ the *comb of* v_i. We have the following property of the construction.

Lemma 1. *In any feasible solution S' of the* Maximum Hidden Set on Polygon with Holes *instance I', at most $2n$ points in S can be placed outside the combs.*

Proof. In each of the n trapezoids $\{v'_{i-1}, v'_i, v_{i,n^2+1}, v_{i,0}\}$ (see Figs. 1 and 2), there can be at most one point, which gives n points in total. Moreover, by our construction any point p in the trapezoid $\{v'_{i-1}, v'_i, m'_i, m'_{i-1}\}$ (not in the holes) can see every point p' in the n-gon $\{v'_0, \ldots, v'_n\}$ except for points p' in any of the holes and (possibly) except for points p' in the triangles $\{v'_{i-1}, m'_{i-1}, e^+_{i-1,i}\}$ and $\{v'_i, m'_i, e^-_{i+1,i}\}$ (see Fig. 2). Therefore, all points in S' that lie in the n-gon $\{v'_0, \ldots, v'_n\}$ must lie in only one of the n polygons $\{e^+_{i-1,i}, m'_{i-1}, m'_i, e^-_{i+1,i}, v'_i, v'_{i-1}\}$. Obviously, at most n points can be hidden in any one of these polygons. □

We have the following observation, which follows directly from the construction:

Observation 1 *Any point in the comb of v_i completely sees the comb of vertex v_j, if (v_i, v_j) is not an edge in the graph. If (v_i, v_j) is an edge in the graph, then no point in the comb of v_i sees any point in the comb of v_j*

Given a feasible solution S' of the Maximum Hidden Set on Polygon with Holes instance I', we obtain a feasible solution S of the Maximum Clique instance I as follows: A vertex $v_i \in V$ is in the solution S, iff at least one point from S' lies in the comb of v_i. To see that S is a feasible solution, assume by contradiction that it is not a feasible solution. Then, there exists a pair of vertices $v_i, v_j \in S$ with no edge between them. But then, there is by construction no hole in the polygon to block the view between the comb of v_i and the comb of v_i.

We need to show that the construction of I' can be done in polynomial time and that a feasible solution can be transformed in polynomial time. There are $2n^2 + 1$ vertices in each of the n combs. We have additional n vertices v'_i. There are 2 holes for each edge in the graph and each hole consists of 3 vertices. Therefore, the polygon P consists of $6m + 2n^3 + 2n$ vertices. It is known in computational geometry that the coordinates of intersection points of lines with rational coefficients can be expressed with polynomial length. All of the points in our construction are of this type. Therefore, the construction is polynomial. The transformation of a feasible solution can obviously be done in polynomial time.

We obtain our inapproximability result by using the technique of gap-preserving reductions (as introduced in [1]), which consists of transforming a promise problem into another promise problem.

Lemma 2. *Let OPT denote the size of an optimum solution of the* Maximum Clique *instance I, let OPT' denote the size of an optimum solution of the* Maximum Hidden Set on Polygon with Holes *instance I', and let $k \leq n$. The following holds: $OPT \geq k \Longrightarrow OPT' \geq n^2 k$*

Proof. If $OPT \geq k$, then there exists a clique in I of size k. We obtain a solution for I' of size $n^2 k$ by simply letting the n^2 vertices $w_{i,l}$ for $l \in \{0, \ldots, n^2\}$ be in the solution if and only if vertex $v_i \in V$ is in the clique. The solution thus obtained for I' is feasible (see Observation 1). □

Lemma 3. *Let OPT denote the size of an optimum solution of the* Maximum Clique *instance I, let OPT' denote the size of an optimum solution of the* Maximum Hidden Set on Polygon with Holes *instance I', let $k \leq n$, and let $\epsilon > 0$. The following holds: $OPT < \frac{k}{n^{1/2-\epsilon}} \Longrightarrow OPT' < \frac{n^2 k}{n^{1/2-\epsilon}} + 2n$*

Proof. We prove the contraposition: $OPT' \geq \frac{n^2 k}{n^{1/2-\epsilon}} + 2n \Longrightarrow OPT \geq \frac{k}{n^{1/2-\epsilon}}$. Suppose we have a solution of I' with $\frac{n^2 k}{n^{1/2-\epsilon}} + 2n$ points. At most $2n$ of the points in the solution can be outside the combs, because of Lemma 1. Therefore, at least $\frac{n^2 k}{n^{1/2-\epsilon}}$ points must be in the combs. From the construction of the combs, it is clear that at most n^2 points can hide in each comb. Therefore, the number of combs that contain at least one point from the solution is at least $\frac{\frac{n^2 k}{n^{1/2-\epsilon}}}{n^2} = \frac{k}{n^{1/2-\epsilon}}$. The transformation of a solution as described above yields a solution of I with at least $\frac{k}{n^{1/2-\epsilon}}$ vertices. □

Lemmas 2 and 3 and the fact that $|I'| \leq 10n^2$ allow us to prove our first main result, using standard concepts of gap-preserving reductions (see [1]). The proof easily carries over to the vertex restricted version of the problem.

Theorem 1. Maximum Hidden Set on Polygon with Holes *and* Maximum Hidden Vertex Set on Polygon with Holes *cannot be approximated by any polynomial time algorithm with an approximation ratio of $\frac{|I'|^{1/6-\gamma}}{4}$, where $|I'|$ is the number of vertices in the polygon, and where $\gamma > 0$, unless $NP = P$.*

3 Inapproximability Results for the Terrain Problems

Theorem 2. *The problems* Maximum Hidden Set on Terrain *and* Maximum Hidden Vertex Set on Terrain *cannot be approximated by any polynomial time algorithm with an approximation ratio of $\frac{|I''|^{1/6-\gamma}}{4}$, where $|I''|$ is the number of vertices in the terrain, and where $\gamma > 0$, unless $NP = P$.*

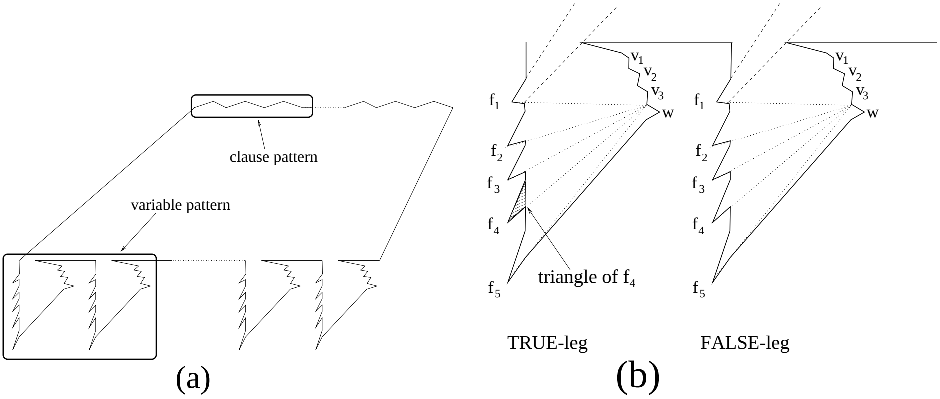

Fig. 4. (a) Schematic construction, (b) Variable pattern

Proof. The proof very closely follows the lines of the proof for the inapproximability of Maximum Hidden (Vertex) Set on Polygon with Holes. We use the same construction, but given the polygon with holes of instance I' we create a terrain (i.e. instance I'') by simply letting all the area outside the polygon (including the holes) have height h and by letting the area in the interior have height 0. We add four vertices to the terrain by introducing a rectangular bounding box around the regular $2n$-gon. This yields a terrain with vertical walls, which can be easily modified to have steep but not vertical walls, as required by the definition of a terrain. Finally, we triangulate the terrain. The terrain thus obtained looks like a canyon of a type that can be found in the south-west of the United States. All proofs work very similar. □

4 Inapproximability Results for the Problems for Polygons Without Holes

We reduce Maximum 5-Occurrence-2-Satisfiability to Maximum Hidden Set on Polygon without Holes to prove the APX-hardness of Maximum Hidden Set on Polygon without Holes. The same reduction will also work for Maximum Hidden Vertex Set on Polygon without Holes with minor modifications. Suppose we are given an instance I of Maximum 5-Occurrence-2-Satisfiability, which consists of n variables $x_0, \ldots, x_{n-1}$ and m clauses $c_0, \ldots, c_{m-1}$. We construct a polygon without holes, i.e. an instance I' of Maximum Hidden Set on Polygon without Holes, which consists of clause patterns and variable patterns, as shown schematically in Fig. 4 (a). The construction uses concepts similar to those used in [9]. The details of the construction are similar to a construction in [6], and will therefore be omitted. It is, however, necessary to introduce the variable pattern. We construct a variable pattern for each variable x_i as indicated in Fig. 4 (b). The cone-like feature drawn with dashed lines simply helps in the construction and is not part of the polygon

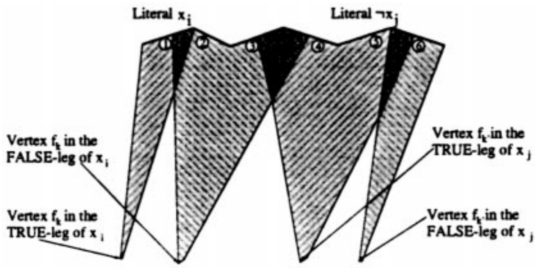

Fig. 5. Clause Pattern with cones

boundary. It represent the link to the clause patterns, as indicated in Fig. 5. Each variable pattern consists of a TRUE- and a FALSE-leg. The reduction has the following properties:

Lemma 4. *If there exists a truth assignment S to the variables of I that satisfies at least $(1-\epsilon)m$ clauses, then there exists a solution S' of I' with $|S'| \geq 10n + 2m + (1-\epsilon)m$.*

Proof. If variable x_i is TRUE in S, then we let the vertices $f_1, \ldots, f_5$ and w of the TRUE-leg of x_i, as well as the vertices v_1, v_2, v_3 and w of the FALSE-leg of x_i be in the solution S'. Vice-versa if x_i is FALSE in S. This gives us $10n$ points in S'. The remaining points for S' are in the clause patterns. Figure 5 shows the clause pattern for a clause $x_i, \neg x_j$[1], together with the cones that link the clause pattern to the corresponding variable patterns. Remember that these cones are not part of the polygon boundary. To understand Fig. 5, assume x_i is assigned the value FALSE and x_j is assigned the value TRUE, i.e., the clause $x_i, \neg x_j$ is not satisfied. Then there is a point in the solution that sits at vertex f_k (for some k) in the FALSE-leg of x_i and a point that sits at vertex f'_k (for some k') in the TRUE-leg of x_j. In this case, we can have only *two* additional points in the solution S' at points 1, 6. In the remaining three cases, where the variables x_i and x_j are assigned truth values such that the clause is satisfied, we can have *three* additional points in S' at 1 – 6. Therefore, we have 2 points from all unsatisfied clauses and 3 points from all satisfied clauses, i.e. $2\epsilon m + 3(1-\epsilon)m$ points that are hidden in the clause patterns. Thus, $|S'| \geq 10n + 2m + (1-\epsilon)m$, as claimed.[2] □

Lemma 5. *If there exists a solution S' of I' with $|S'| \geq 10n + 3m - (\epsilon + \gamma)m$, then there exists a variable assignment S of I that satisfies at least $(1 - \epsilon - \gamma)m$ clauses.*

Proof. For any solution S', we can assume that in each leg of each variable pattern, all points in S' are either in the triangles of vertices $f_1, \ldots, f_5$ and w,

[1] The proofs work accordingly for other types of clauses, such as x_i, x_j.

[2] Note that a point at some vertex f_k actually sees a slightly larger cone than indicated in Fig. 5. This problem can be dealt with by making the triangle of f_k very small. Corresponding methods are used to solve similar problems in [6] and [4].

or in the triangles of vertices v_1, v_2, v_3, and w (see Fig. 4 (b) for the definition of these triangles), since any point in any triangle of $f_1, \ldots, f_5$ sees the triangles of v_1, v_2, v_3 completely and any single point in the leg outside these triangles would see almost all (at least 3) of these triangles, and we could obtain better solutions easily. We transform the solution S' (with $|S'| \geq 10n + 3m - (\epsilon + \gamma)m$) in such a way that it remains feasible and that its size (i.e. the number of hidden points) does not decrease. This is done with an enumeration of all possible cases, i.e. we show how to transform the solution if there is a point in 3, 4, or 5 of the triangles of the points $f_1, \ldots, f_5$ in the TRUE-leg and the FALSE-leg of a variable pattern. The transformation is such that at the end, we have for each variable pattern the six points at $f_1, \ldots, f_5$, and w from one leg in the solution and the 4 points v_1, v_2, v_3, and w from the other leg. Thus, we can easily obtain a truth assignment for the variables by letting variable x_i be TRUE iff the six points at $f_1, \ldots, f_5$, and w from the TRUE-leg are in the solution. The transformed solution S' consists of at least $10n + 3m - (\epsilon + \gamma)m$ points, $10n$ of which lie in the variable patterns. At most 3 points can lie in each clause pattern. If 3 points lie in a clause pattern, then this clause is satisfied. Therefore, if 2 points lie in each clause pattern, there are still at least $(1 - \epsilon - \gamma)m$ additional points in S'. These must lie in clause patterns as well. Therefore, at least $(1 - \epsilon - \gamma)m$ clauses are satisfied. □

Lemmas 4 and 5 show how to transform two promise problems into one another. By using standard concepts of gap-preserving reductions and by introducing some minor modification for the vertex-restricted problem, we obtain:

Theorem 3. Maximum Hidden (Vertex) Set on Polygon without Holes *is APX-hard, i.e. there exists a constant $\delta > 0$ for each of the two problems such that no polynomial time approximation algorithm for the problem can achieve an approximation ratio of $1 + \delta$, unless $P = NP$.*

5 Conclusion

We have shown that the problems Maximum Hidden (Vertex) Set on Polygon with Holes and Maximum Hidden (Vertex) Set on Terrain are almost as hard to approximate as Maximum Clique. We could prove for all these problems an inapproximability ratio of $O(|I'|^{1/3-\gamma})$, but under the assumption that $coR \neq NP$, using the stronger inapproximability result for Maximum Clique from [7]. Furthermore, we have shown that Maximum Hidden (Vertex) Set on Polygon without Holes is APX-hard. Note that an approximation algorithm for all considered problems that simply returns a single vertex achieves an approximation ratio of n. Note that our proofs can easily be modified to work as well for polygons or terrains, where no three vertices are allowed to be collinear. We have classified the problems Maximum Hidden (Vertex) Set on Polygon with Holes and Maximum Hidden (Vertex) Set on Terrain to belong to the class of problems inapproximable with an approximation ratio of n^{ϵ} for some $\epsilon > 0$, as defined in [1]. The APX-hardness results for the problems for polygons without holes, however, do not precisely characterize the approximability characteristics of these problem. The gap between the best (known) achievable approximation ratio (which is n) and the best inapproximability ratio is still very large for these problems and should be closed in future research. As for other future work, we plan to consider several variations of the problems presented. For example, we plan to try to hide non-zero-dimensional objects.

References

1. S. Arora and C. Lund; *Hardness of Approximations*; in: Approximation Algorithms for NP-Hard Problems (ed. Dorit Hochbaum), PWS Publishing Company, pp. 399-446, 1996.
2. P. Bose, T. Shermer, G. Toussaint, and B. Zhu; *Guarding Polyhedral Terrains*; Computational Geometry 7, pp. 173-185, Elsevier Science B. V., 1997.
3. P. Crescenzi, V. Kann; *A Compendium of NP Optimization Problems*; in the book by G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, M. Protasi, *Complexity and Approximation. Combinatorial Optimization Problems and their Approximability Properties*, Springer-Verlag, Berlin, 1999; also available in an online-version at: http://www.nada.kth.se/theory/compendium/compendium.html.
4. S. Eidenbenz, C. Stamm, and P. Widmayer; *Inapproximability of some Art Gallery Problems*; Proc. 10th Canadian Conf. Computational Geometry, pp. 64-65, 1998.
5. S. Eidenbenz, C. Stamm, and P. Widmayer; *Positioning Guards at Fixed Height above a Terrain – An Optimum Inapproximability Result*; Lecture Notes in Computer Science, Vol. 1461 (ESA'98), p. 187-198, 1998.
6. S. Eidenbenz; *Inapproximability Results for Guarding Polygons without Holes*; Lecture Notes in Computer Science, Vol. 1533 (ISAAC'98), p. 427-436, 1998.
7. J. Hastad; *Clique is Hard to Approximate within* $n^{1-\epsilon}$; Proc. of the Symposium on Foundations of Computer Science, 1996.
8. M. van Kreveld; *Digital Elevation Models and TIN Algorithms*; in Algorithmic Foundations of Geographic Information Systems (ed. van Kreveld et al.), LNCS tutorial vol. 1340, pp. 37-78, Springer, 1997.
9. D. T. Lee and A. K. Lin; *Computational Complexity of Art Gallery Problems*; IEEE Trans. Info. Th, pp. 276-282, IT-32, 1986.
10. J. O'Rourke; *Art Gallery Theorems and Algorithms*; Oxford University Press, New York, 1987.
11. C.H. Papadimitriou; *Computational Complexity*; Addison-Wesley Publishing Company, 1994.
12. T. Shermer; *Hiding People in Polygons*; Computing 42, pp. 109 - 131, 1989.
13. T. Shermer; *Recent results in Art Galleries*; Proc. of the IEEE, 1992;
14. J. Urrutia; *Art Gallery and Illumination Problems*; in Handbook on Computational Geometry, edited by J.-R. Sack and J. Urrutia, 1998.
15. B. Zhu; *Computing the Shortest Watchtower of a Polyhedral Terrain in* $O(n \log n)$ *Time*; in Computational Geometry 8, pp. 181-193, Elsevier Science B.V., 1997.